arXiv:1411.0591v1 [cond-mat.stat-mech] 3 Nov 2014

Bayesian feature selection with strongly-regularizing priors maps to the Ising Model

Charles K. Fisher and Pankaj Mehta
Dept. of Physics, Boston University, Boston, MA 02215
(Dated: November 4, 2014)

Identifying small subsets of features that are relevant for prediction and/or classification tasks is a central problem in machine learning and statistics. The feature selection task is especially important, and computationally difficult, for modern datasets where the number of features can be comparable to, or even exceed, the number of samples. Here, we show that feature selection with Bayesian inference takes a universal form and reduces to calculating the magnetizations of an Ising model, under some mild conditions. Our results exploit the observation that the evidence takes a universal form for strongly-regularizing priors — priors that have a large effect on the posterior probability even in the infinite data limit. We derive explicit expressions for feature selection for generalized linear models, a large class of statistical techniques that include linear and logistic regression. We illustrate the power of our approach by analyzing feature selection in a logistic regression-based classifier trained to distinguish between the letters B and D in the notMNIST dataset.

Modern technological advances have fueled the growth of large datasets where thousands, or even millions, of features can be measured simultaneously. Dealing with such a large number of features is often impractical however, especially when the sample size is limited. Feature selection is one promising approach for dealing with such complex datasets [1, 2]. The goal of feature selection is to identify a subset of features that are relevant for statistical tasks such as prediction or classification. Feature selection is a difficult problem because features are often redundant and strongly correlated with each other. Furthermore, the relevance of a feature depends not only on the dataset being analyzed, but also the statistical techniques being employed. While powerful methods for feature selection exist for a handful of techniques such as linear regression [3, 4], there is no unified framework for performing feature selection in a computationally tractable manner. To address this shortcoming, we present a unified approach for Bayesian feature selection that is applicable to a large class of commonly used statistical models.

In principle, Bayesian inference provides a unified framework for feature selection [2, 5]. Unfortunately, Bayesian methods often require extensive Monte Carlo simulations that become computationally intractable for very high-dimensional problems. In the Bayesian framework, a statistical model is defined by its likelihood function, $p_x(y|\boldsymbol{\theta})$, which describes the observed data, y, as a function of some features, $\boldsymbol{x}$, and model parameters, $\boldsymbol{\theta}$. This likelihood function is supplemented by a prior, $p(\boldsymbol{\theta})$, that encodes our belief about the model parameters in the absence of data. Using Bayes rule, one can define the posterior probability, $p_x(\boldsymbol{\theta}|y) \propto p_x(y|\boldsymbol{\theta})p(\boldsymbol{\theta})$, that describes our belief about the parameter $\boldsymbol{\theta}$ given the observed data y. Notice, that for a flat (constant) prior, maximizing the posterior distribution is equivalent to maximizing the likelihood and Bayesian inference reduces to the usual maximum likelihood framework. More generally, the log of the prior can be interpreted as a penalty function that "regularizes" the model parameters. For example, the choice of a Gaussian and Laplace prior correspond to imposing a $L2$ and $L1$ penalty on model parameters, respectively [1].

Since the ultimate goal of feature selection is to identify a subset of features, it is helpful to introduce a set of indicator variables, $\boldsymbol{s}$, that indicates whether a feature is included in the statistical model, with $s_i = 1$ if feature i is included and $s_i = -1$ if it is not. Assuming that we have no *a priori* information about which variables are relevant, the posterior distribution for $\boldsymbol{s}$ can be computed as [6]:

$$p_x(\boldsymbol{s}|y) \propto p_x(y|\boldsymbol{s}). \tag{1}$$

where the 'evidence' is given by:

$$p_x(y|\boldsymbol{s}) = \int d\boldsymbol{\theta}\, p_x(y|\boldsymbol{\theta}, \boldsymbol{s}) p(\boldsymbol{\theta}|\boldsymbol{s}) \tag{2}$$

Feature selection can be performed by choosing features with large posterior probabilities.

Our central result is that the logarithm of the evidence maps to the energy of an Ising model for a large class of priors we term "strongly-regularizing priors". Thus, computing the marginal posterior probabilities for the relevance of each feature reduces to computing the magnetizations of an Ising model. The key property of strongly-regularizing priors is that they affect the posterior probability even when the sample size goes to infinity [7]. This should be contrasted with the usual assumption of Bayesian inference the effect of the prior on posterior probabilities diminishes inversely with the number of samples [2]. Strongly-regularizing priors are especially useful for feature selection problems where the number of potential features is greater than, or comparable to, the number of data points. Surprisingly, our results do not depend on the specific choice of prior or likelihood function, under some mild conditions, suggesting that Bayesian feature selection has universal properties for strongly regularizing priors.

Practically, our Bayesian Ising Approximation (BIA) provides an efficient method for feature selection with many commonly employed statistical procedures such as Generalized Linear Models (GLMs). We envision using the BIA as part of a two-stage procedure where the BIA is applied to rapidly screen irrelevant variables, i.e. those features that have low rank in posterior probability, before applying a more computationally intensive cross validation procedure to infer the model parameters with the remaining features.

GENERAL FORMULAS

For concreteness, consider a series of n independent observations of some dependent variable y, i.e. $\boldsymbol{y} = (y_1, \ldots, y_n)$, that we want to explain using a set of p potential features $\boldsymbol{x} = (x_1, \ldots, x_p)$. Furthermore, assume that to each of these potential features, x_i, we associate a model parameter θ_i. Let S denote an index set specifying the positions j for which $s_j = +1$. The cardinality of S, i.e. the number of indicator variables with $s_j = +1$, is denoted M. Also, we define a vector $\hat{\boldsymbol{\theta}}$ of length M that contains the model parameters corresponding to the active features, S. With these definitions, we define the log-likelihood for the observed data as a function of the active features as $\mathcal{L}_X(\boldsymbol{y}|\hat{\boldsymbol{\theta}}) = -\log p_X(\boldsymbol{y}|\hat{\boldsymbol{\theta}})$. Throughout, we assume that the log-likelihood is twice differentiable.

In addition to the log-likelihood, we need to specify prior distributions for the parameters. Here, we will work with factorized priors of the form

$$P(\hat{\boldsymbol{\theta}}) = \prod_{j \in S} \frac{1}{Z(\lambda)} e^{-\lambda f(\hat{\theta}_j)}, \tag{3}$$

where λ is a parameter that controls the strength of the regularization, $f(\theta_j)$ is a twice differentiable convex function minimized at the point $\bar{\theta}_j$ with $f(\bar{\theta}_j) = 0$, and $Z(\lambda)$ is the normalization constant. As the function is convex, we are assuming that the second derivative evaluated at $\bar{\theta}_j$, denoted $\partial^2 f$, is positive. Many commonly used priors take this form including Gaussian and hyperbolic priors. Plugging these expressions into (2), yields an expression for the evidence when only the subset of features, S, are included:

$$p_X(\boldsymbol{y}|S) = Z^{-M}(\lambda) \int d\hat{\boldsymbol{\theta}} e^{\mathcal{L}_X(\boldsymbol{y}|\hat{\boldsymbol{\theta}}) - \lambda \sum_{j \in S} f(\hat{\theta}_j)} \tag{4}$$

By definition, this integral is dominated by the second term for strongly regularizing priors. Since the log-likelihood is extensive in the number of data points, n, this generally requires that the regularization strength be much larger than the number of data points, $\lambda \gg n$. For such strongly regularizing priors we can perform a saddle-point approximation for the evidence. This yields, up to an irrelevant constant,

$$\begin{aligned} \mathcal{L}_X(\boldsymbol{y}|S) &= \log p_X(\boldsymbol{y}|S) \\ &= -\frac{1}{2}\log|I - \Lambda^{-1}H| + \frac{\Lambda^{-1}}{2}\boldsymbol{b}'(I - \Lambda^{-1}H)^{-1}\boldsymbol{b} \end{aligned} \tag{5}$$

where $\Lambda = \lambda(\partial^2 f)$ is a "renormalized" regularization strength,

$$\boldsymbol{b} = \nabla \mathcal{L}_X(\boldsymbol{y}|\hat{\boldsymbol{\theta}}) \tag{6}$$

is the gradient of the log-likelihood, and

$$H_{ij} = \frac{\partial^2 \mathcal{L}_X(\boldsymbol{y}|\hat{\boldsymbol{\theta}})}{\partial \hat{\theta}_i \partial \hat{\theta}_j} \tag{7}$$

is the Hessian of the log-likelihood, with all derivatives evaluated at $\hat{\theta}_j = \bar{\theta}_j$. We emphasize that the large parameter in our saddle-point approximation is the regularization strength $\Lambda \gg n$. This differs from previous statistical physics approaches to Bayesian feature selection that commonly assume that $n \gg \Lambda$ and hence perform a saddle point in n [8–10]. The fundamental reason for this difference is that we work in the strongly-regularizing regime where the prior is assumed to be important even in the infinite data limit.

Since for strongly-regularizing priors, Λ is the largest scale in the problem, we can expand the log in (5) in a power series expansion in $\epsilon = \Lambda^{-1}$ to order $O(\epsilon^3)$ to get:

$$\mathcal{L}_X(\boldsymbol{y}|S) \simeq \frac{\epsilon}{2}(\mathrm{Tr}[H] + \boldsymbol{b}'\boldsymbol{b}) + \frac{\epsilon^2}{2}(\frac{1}{2}\mathrm{Tr}[H^2] + \boldsymbol{b}'H\boldsymbol{b}) \tag{8}$$

One of the most striking aspects of this expression is that it is independent of the detailed form of the prior function, $f(\theta)$. All that is required is that the prior is a twice differential convex function with a global minimum at $f(\bar{\theta}_j) = 0$. Different choices of prior simply "renormalize" the effective regularization strength Λ.

Rewriting this expression in terms of the spin variables, $\boldsymbol{s}$, we see that the evidence takes the Ising form with:

$$\mathcal{L}_X(\boldsymbol{y}|\boldsymbol{s}) = \epsilon \left(\sum_i s_i h_i + \sum_{ij} J_{ij} s_i s_j \right) \tag{9}$$

and couplings given by:

$$h_i = \frac{1}{4}(H_{ii} + b_i^2) + 2\sum_j J_{ij} \tag{10}$$

$$J_{ij} = \frac{\epsilon}{16}(H_{ij}^2 + 2b_i H_{ij} b_j) \tag{11}$$

Notice that the couplings, which are proportional to the small parameter ϵ, are weak. According to Bayes rule, $\mathcal{L}(\boldsymbol{y}|\boldsymbol{s}) = \mathcal{L}(\boldsymbol{s}|\boldsymbol{y}) + constant$, so the posterior probability of a set of features $\boldsymbol{s}$ is described by an Ising model of the form:

$$p(\boldsymbol{s}|\boldsymbol{y}) = \mathcal{Z}^{-1} e^{\sum_i s_i h_i + \sum_{ij} s_i s_j J_{ij}} \tag{12}$$

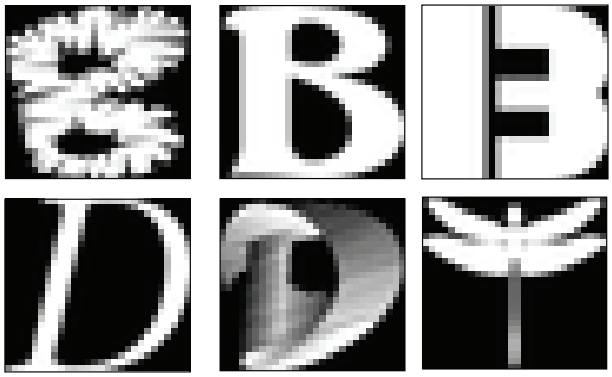

FIG. 1. **Three randomly chosen examples of Bs (top) and Ds (bottom) from the notMNIST dataset.** Each letter in the notMNIST dataset is represented by a 28 by 28 pixel grayscale image.

where $\mathcal{Z}$ is the partition function that normalizes the probability distribution. We term this the Bayesian Ising Approximation (BIA) [7]. Finally, for future use, it is useful to define a scale Λ^* for which the Ising approximation breaks down. This scale Λ^* can be computed by requiring the that the power series expansion used to derive (8) converge [7].

We demonstrated the utility of the BIA for feature selection in the specific context of linear regression in a recent paper, and we can import that machinery here [7]. It is useful to explicitly indicate the dependence of various expression on the regularization strength Λ. We want to compute the marginal posterior probability that a feature, j, is relevant:

$$p_\Lambda(s_j = 1|\mathbf{y}) \simeq (1 + m_j(\Lambda))/2 \tag{13}$$

where we have defined the magnetizations $m_j(\Lambda) = \langle s_j \rangle$. While there are many techniques for calculating the magnetizations of an Ising model, we focus on the mean field approximation which leads to a self-consistent equation [11]:

$$m_i(\Lambda) = \tanh\left[\frac{n^2}{4\Lambda}\left(h_i(\Lambda) + \frac{1}{2}\sum_{j\neq i} J_{ij}(\Lambda) m_j(\Lambda)\right)\right]$$

Our expressions depend on a free parameter (Λ) that determines the strength of the prior distribution. As it is usually difficult, in practice, to choose a specific value of Λ ahead of time it is often helpful to compute the feature selection path; i.e. to compute $m_j(\Lambda)$ over a wide range of Λ's. Indeed, computing the variable selection path is a common practice when applying other feature selection techniques such as LASSO regression [3]. To obtain the mean field variable selection path as a function of $\epsilon = 1/\Lambda$, we notice that $\lim_{\epsilon\to 0} m_j(\epsilon) = 0$ and so define the recursive formula:

$$m_i(\epsilon + \delta\epsilon) = \tanh\left[\frac{(\epsilon+\delta\epsilon)n^2}{4}\left(h_i\left(\epsilon+\delta\epsilon\right) + \frac{1}{2}\sum_{j\neq i} J_{ij}\left(\epsilon+\delta\epsilon\right) m_j\left(\epsilon\right)\right)\right]$$

with a small step size $\delta\epsilon \ll 1/\Lambda^*$. We have set $\delta\epsilon = 0.05/\Lambda^*$ in all of the examples presented below.

EXAMPLES

Logistic regression is a commonly used statistical method for modeling categorical data [1]. To simplify notation, it is useful also useful define an extra feature variable $x_0 = 1$ which is always equal to 1 and a $p+1$-dimensional vector of feature, $\boldsymbol{x} = (x_0, x_1, \ldots, x_p)$ with corresponding parameters, $\boldsymbol{\theta} = (\theta_0, \theta_1, \ldots, \theta_p)$. In terms of these vectors, the likelihood function for logistic regression takes the compact form

$$P_{\boldsymbol{x}}(y = 1|\boldsymbol{\theta}) = 1 - P_{\boldsymbol{x}}(y = 1|\boldsymbol{\theta}) = \frac{e^{\boldsymbol{\theta}'\boldsymbol{x}}}{1 + e^{\boldsymbol{\theta}'\boldsymbol{x}}}, \tag{14}$$

where $\boldsymbol{\theta}'$ is the transpose of $\boldsymbol{\theta}$. If we have n independent observations of the data (labeled by index $l = 1 \ldots n$), the log-likelihood can be written as:

$$\mathcal{L}_X(\boldsymbol{y}|\boldsymbol{\theta}) = \sum_l y_l(\boldsymbol{\theta}'\boldsymbol{x}^l) - \log\left(1 + e^{\boldsymbol{\theta}'\boldsymbol{x}^l}\right). \tag{15}$$

We supplement this likelihood function with an L^2 norm on the parameters of the form:

$$p(\boldsymbol{\theta}) = \prod_{j=1}^{p} \sqrt{\frac{\lambda}{2\pi}} e^{-\lambda(\theta_j - \bar{\theta}_j)^2/2} \tag{16}$$

with $\bar{\boldsymbol{\theta}} = (\bar{\theta}_0, 0, \ldots, 0)$ and where $\bar{\theta}_0$ is chosen to match the observed probability that $y = 1$ in the data,

$$P_{obs}(y). = \frac{e^{\bar{\theta}_0}}{1 + e^{\bar{\theta}_0}} \tag{17}$$

Using these expressions, we can calculate the gradient and the Hessian of the log-likelihood:

$$b_i \equiv \frac{\partial \mathcal{L}_X(\boldsymbol{y}|\boldsymbol{\theta})}{\partial \theta_i}|_{\bar{\boldsymbol{\theta}}} = \sum_l x_i^l \left(y^l - P_{obs}(y)\right) \tag{18}$$

and

$$H_{ij} = -\left(P_{obs}(y)(1 - P_{obs}(y))\right. \cdot \sum_l x_i^l x_j^l \tag{19}$$

Plugging in these into (11) yields the couplings for the Ising model in (12).

Notice that the gradient, b_i, is proportional to the correlations between the y and x_i. Furthermore, except for

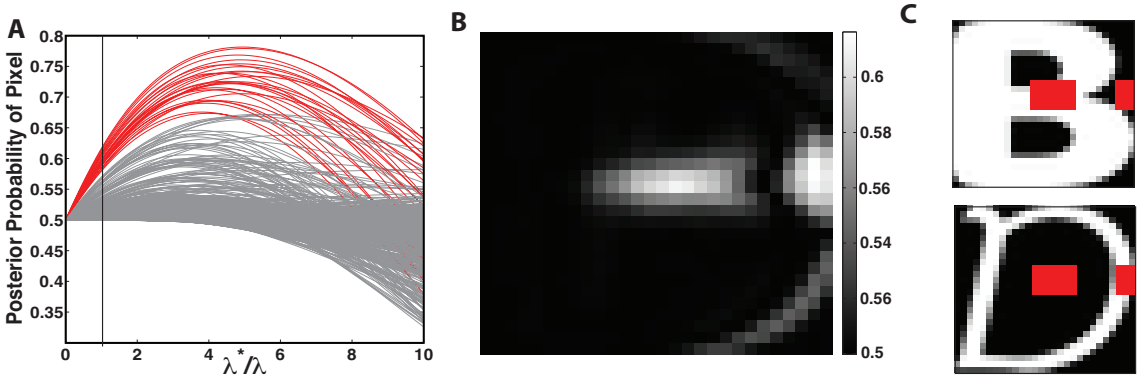

FIG. 2. **Identifying the most informative pixels for distinguishing Bs and Ds**. **A** We used the Bayesian Ising Approximation (BIA) to perform feature selection for a logistic regression based classifier designed to classify Bs and Ds. The model was trained 500 using randomly chosen examples of each letter from the notMNIST dataset. The resulting feature selection path is shown as a function of the regularization strength, λ. Each line corresponds to one of 784 pixels. We expect the BIA to breakdown when $\lambda < \lambda^*$. Pixels in red have the highest posterior probability and hence can be used for distinguishing between Bs and Ds. **B** Posterior probabilities of all 784 pixels when $\Lambda/\Lambda^* = 1$. **C** Visualization of the most informative pixels labeled in red in part A.

a multiplicative constant reflecting the variance of y, H_{ij} is just correlation between x_i and x_j. Thus, as in linear regression [7], the coefficients of the Ising model are related to the correlations between variables and/or the data. In fact, it is easy to show this is the case for all Generalized Linear Models (see Appendix).

To illustrate our approach, we used the BIA for a logistic regression-based classifier designed to classify Bs and Ds in the notMNIST dataset. The notMNIST dataset consists of diverse images of the letters A-J composed from publicly available computer fonts. Each letter is represented as a 28 by 28 grayscale image where pixel intensities vary between 0 and 255. Figure 1 shows three randomly chosen examples of the letters B and D from the notMNIST dataset. The BIA was performed using 500 randomly chosen examples each letter. Notice that the number of examples (500 or each letter) is comparable to the number of pixels (784) in an image, suggesting that strongly-regularizing the problem is appropriate.

Using the expressions above, we calculated the couplings for the Ising model describing our logistic-regression based classifier and calculated the feature selection path as a function of Λ/Λ^* using the mean field approximation. As in linear regression, we used $\Lambda^* = n(1 + pr)$, where $n = 1000$ is the number of samples, $p = 784$ is the number of potential features, and r is the root mean squared correlation between pixels (see Figure 2A). Figure 2B shows the posterior probability of all 784 pixels when $\Lambda = \Lambda^*$. To better visualize this, we have labeled the pixels with the highest posterior probabilities in red in the feature selection path in Fig 2A and in the sample images shown Figure 2C. The results agree well with our basic intuition about which pixels are important for distinguishing the letters B and D.

Conclusion

In conclusion, we have presented a general framework for Bayesian Feature Selection in a large class of statistical models. In contrast to previous Bayesian approaches that assume that the effect of the prior vanishes inversely with the amount of data, we have used strongly-regularizing priors that have large effects on the posterior distribution even in the infinite data limit. We have shown that in the strongly-regularized limit, Bayesian feature selection takes the universal form of an Ising model. Thus, the marginal posterior probabilities that of each feature is relevant can be efficiently computed using a mean-field approximation. Furthermore, for Generalized Linear Models we have shown the coefficients of the Ising model can be calculated directly from correlations between the data and features.

Surprisingly, aside from some mild regularity conditions, our approach is independent of the choice of prior or likelihood. This suggests that it maybe possible to obtain general results about strongly-regularizing priors. It will be interesting to further explore Bayesian inference in this new limit. Our approach also gives a practical algorithm for quickly performing feature selection in many commonly employed statistical and machine learning approaches The methods outlined here are especially well suited for modern data sets where the number of potential features can vastly exceed the number of independent samples. We envision using the Bayesian feature selection algorithm outlined here as part of a two stage procedure. One can use the BIA to rapidly screen irrelevant variables and reduce the complexity of the dataset before applying a more comprehensive cross-validation procedure. More generally, it will be interesting to further develop statisti-

cal physics based approaches for the analysis of complex data.

Appedix: Formulas for Generalized Linear Models

The gradient. $\boldsymbol{b}$, and Hessian, H, of the log-likelihood have particularly simple definitions for Generalized Linear Models (GLMs) which extend the exponential family of distributions. In the exponential family, we can write the distribution in the form

$$p(y|\theta) = g(y)e^{\boldsymbol{\eta}(\theta)\boldsymbol{T}(y)+F(\eta)} \tag{20}$$

where $\boldsymbol{T}(y)$ is a vector of sufficient statistics, and the $\boldsymbol{\eta}$ are called natural parameters. Notice, that for these distributions, we have that

$$\frac{\partial F}{\partial \eta_i} = -\langle T_i(y)\rangle \tag{21}$$

and

$$\frac{\partial^2 F}{\partial \eta_i \partial \eta_j} = \mathrm{Cov}[T_i(y), T_j(y)] \tag{22}$$

where Cov denotes the covariance (connected correlation function).

In a GLM, we restrict ourselves to distribution to scalar quantities where $T(y) = y$ and say that $\eta = \boldsymbol{\theta}'\boldsymbol{x}$. Then, we can write the likelihood as

$$p(y|\boldsymbol{\theta}, \boldsymbol{x}) = g(y)e^{(\boldsymbol{\theta}'\boldsymbol{x})y+F(\boldsymbol{\theta}'\boldsymbol{x})} \tag{23}$$

If we have n independent data points with $l = 1 \ldots n$ then we can write the log-likelihood for such a distribution as

$$\mathcal{L}_X(\boldsymbol{y}|\boldsymbol{\theta}) = \sum_l \log g(y^l) + (\boldsymbol{\theta}'\boldsymbol{x}^l)y^l + F(\boldsymbol{\theta}'\boldsymbol{x}^l) \tag{24}$$

Using the expressions above for the exponential family and (6) we have that

$$b_i = \frac{\partial \mathcal{L}_X(\boldsymbol{y}|\boldsymbol{\theta})}{\partial \theta_i}|_{\boldsymbol{\theta}=\bar{\boldsymbol{\theta}}} = \sum_l y^l x_i^l - x_i^l \langle y\rangle_{\bar{\boldsymbol{\theta}}} \tag{25}$$

where $\langle y\rangle_{\bar{\boldsymbol{\theta}}}$ is the expectation value of y for choice of parameter $\boldsymbol{\theta} = \bar{\boldsymbol{\theta}}$. If we choose $\bar{\theta}$ to reproduce the empirical probability we get:

$$b_i \approx n\mathrm{Cov}[y, x_i] \tag{26}$$

Moreover, the entries of the Hessian are given by:

$$H_{ij} = \frac{\partial^2 \mathcal{L}_X(\boldsymbol{y}|\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} = -\sum_l x_i^l x_j^l \mathrm{Var}[y]_{\boldsymbol{\theta}=\bar{\boldsymbol{\theta}}} \tag{27}$$

If we consider standardized variables, $\boldsymbol{x}$, then we can write:

$$H_{ij} \approx -n\mathrm{Cov}[x_i, x_j]\mathrm{Var}[y] \tag{28}$$

[1] C. M. Bishop *et al.*, *Pattern recognition and machine learning*, Vol. 1 (springer New York, 2006).

[2] D. J. Mackay, Information Theory, Inference and Learning Algorithms, by David JC MacKay, pp. 640. ISBN 0521642981. Cambridge, UK: Cambridge University Press, October 2003. **1** (2003).

[3] R. Tibshirani, Journal of the Royal Statistical Society. Series B (Methodological) , 267 (1996).

[4] H. Zou and T. Hastie, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **67**, 301 (2005).

[5] A. O'Hagan, J. Forster, and M. G. Kendall, *Bayesian inference* (Arnold London, 2004).

[6] D. J. MacKay, Neural computation **4**, 415 (1992).

[7] C. K. Fisher and P. Mehta, arXiv preprint arXiv:1407.8187 (2014).

[8] J. B. Kinney and G. S. Atwal, Neural computation **26**, 637 (2014).

[9] V. Balasubramanian, Neural computation **9**, 349 (1997).

[10] I. Nemenman and W. Bialek, Physical Review E **65**, 026137 (2002).

[11] M. Opper and O. Winther, Advanced mean field methods: theory and practice , 7 (2001).